

Multipoint Sensing by Intelligent Collectives

Loren P. Clare
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
(818) 354-1650
Loren.P.Clare@jpl.nasa.gov

Abstract—Many remote sensing applications require that multiple sensors collect data simultaneously at spatially distributed locations and their information combined in order to characterize the phenomena of interest. Several basic classes of such multipoint measurement systems may be identified. For each, centralized methods exist for combining the raw data from the various sensors. However, recent advancements have given rise to small, integrated nodes comprised of one or more miniaturized sensors, processor, wireless communications capability and power supplies. Collections of these may be deployed and self-organized into intelligent sensor networks capable of performing cooperative signal processing locally, thereby providing substantial benefits.

1. INTRODUCTION

Multipoint measurement systems utilize multiple sensors that are spatially dispersed and simultaneously collect sensor signal data. By aggregating the information from the various sensors, certain characteristics can be determined that may not be possible otherwise. It appears that applications having these qualities may be grouped into a few basic classes. While conventional implementations transport the raw sensor signal data to an end user for centralized processing, recent advancements in microsensor networks [1] offer alternatives that enable local processing of the data. Such an approach can greatly reduce the communications costs, since the processed data generally has much less volume and therefore much lower bandwidth is needed for long haul transport to the end user(s).

Furthermore, depending on the class of multipoint measurement application, the associated algorithms for combining the different sensor signal streams may be broken into distributed processing components that can be performed by different nodes in parallel. In these cases, additional benefits may arise by spreading the processing and communications workload over a collective of cooperating nodes [2].

In the next section, we present three classes of multipoint measurement applications, with illustrative instances. Section 3 describes the intelligent collective concept, in which homogeneous wireless sensor nodes operate cohesively in a network. Section 4 provides some example cases, including associated signal processing techniques, and suggests technology extensions toward achieving the concept. The final section provides conclusions.

2. MULTIPOINT MEASUREMENT CLASSES

Use of multiple sensors to characterize a target phenomenon is often performed by fusing data from sensors having different modalities, such as seismic and acoustic. It may be that these sensors may operate better when deployed at different locations, e.g., due to their having different sensing ranges. However, in this paper, we *assume that the sensors are of similar nature*, having only a single sensing modality. We ask the question: why would one need to use many spatially distributed sensors of the same type? One reason might be to provide fault tolerance through simple redundancy. However, of greater interest is when characteristics of the target

phenomenon can *only* be sensed via multipoint measurement.

We offer below three classes of applications for which multipoint sensing is fundamental in this sense. While these appear to cover a wide range of applications, no claim is made that all multipoint sensing applications must fall into one of these classes. In addition, there can be overlap between these class definitions. Nevertheless, it is hoped that these will provide qualitative distinctions that will aid in identifying supporting general algorithms, services and protocols that will accelerate technological development of intelligent sensor networks. The classes are:

Pixellation (or voxellation) of space. In this class, the target phenomenon to be sensed has a spatial extent that is considerably larger than the range of an individual sensor. Each sensor is capable of capturing a piece of the overall “picture” of the target which is determined by stitching together the neighboring sensor nodes (much like individual pixels that make up an image). For example, the structure (gradients) of a bio-chemical cloud might be derived by combining the results of many spatially dispersed sensors. An example space science application is the proposed Magnetospheric Constellation Mission [3], which calls for up to 100 nanosatellites in various Earth orbits ranging to 60 Earth radii, each of which has a magnetometer, a electron detector and an ion detector, with the objective of characterizing the Earth’s magnetosphere.

A basic aspect of the “Pixellation” class is that the sensed object is dynamic. Otherwise, a single sensor could be used by moving it through the space of interest, recording the results, and subsequently combining the data set into a complete characterization.

It is noted that many short-range sensors may need to be deployed simply to provide adequate coverage of a large area/volume. Even though the target may be a point source viewed completely by exactly one sensor, this may be considered a “Pixellation” case because of the dynamic

appearance of the target and its limited signature. In particular, the target track would be a space-time object (curve in the picture) determined by the overall sensor array.

Beamformation. Another class of multipoint measurement nominally involves a single point source that emits a signal that traverses different paths as it is sensed by different individual sensors. Each different path will generally produce different effects, such as propagation delay and attenuation, as well as add noise that is uncorrelated (at least to some degree). By combining the separate sensor signals, one is able to (1) increase the sensor signal to noise ratio, and (2) localize the position (velocity, etc.) of the point source.

Of course, more generally there can be numerous point sources present simultaneously. In this case, beamformers are mathematically created to isolate the different point sources, and may also be formed to “null” point sources that would otherwise confound characterization of the desired target.

Included in this class are interferometers, in which a narrowband source is characterized by constructive and destructive wave combining of sensor signals from sensors deployed with a precise spatial arrangement. An example space science application is the Terrestrial Planet Finder [4] mission. Techniques for beamformation and localization using wideband signals from randomly deployed sensors may be found in [5].

Tomography. The third class involves sensors each of which captures a lower-dimensional projection of the sensed object from a different perspective, such that by combining the different “images” the higher-dimensional object may be reconstructed. As was for the “Pixellation” case, the target object has dynamics that prevents use of an individual sensor that can be moved around the target (such as the case of CAT scans). An example Earth science application is the Leonardo mission [6], which will use multiangular measurement methods by means of multiple

satellites that each image a common target location from a different perspective.

3. INTELLIGENT SENSOR COLLECTIVE CONCEPT

An intelligent collective that performs multipoint measurements is based on a group of nodes, each of which consists of the same type sensor, a processor (for signal processing and communications protocol execution), a radio, and power supply (e.g., battery). This collection of nodes self-organizes into a cohesive sensor network, and has the following properties:

Signal processing is performed locally. Sensing is very often performed remotely from the end user. Rather than hauling back high bandwidth raw sensor signals from all of the sensor nodes to the end user for processing, there is sufficient intelligence within the collective to perform in situ processing. This will generate a result having a higher level of abstraction that may be represented by a far smaller amount of bits, thereby significantly saving the long haul communications costs. In particular, it may be possible to operate in an event-based fashion, so that long haul communications is only used when the event(s) of interest has (have) actually occurred (other than occasional “heartbeat” indicators of live status). Processing may be limited to that subset of signals associated with sensors that determine they are currently sensing the object(s) of interest. Also, inconsistent readings by a faulty node may be deduced quickly.

If possible, distribute the signal processing. Many times, the sensor signal combining operations may be divided into many subtasks, each of which only depends on a local set of data. Such architectures are desirable for implementation, even if some degradation in pure quality of the result is sacrificed by being outweighed by the ensuing benefits. Distributed operation may significantly reduce the local communications load, particularly if communications protocol primitives that match the distributed processing are provided. Furthermore, the processing load is more evenly balanced, spreading the energy use over a wider

set of nodes as well as speeding the computation through parallelism.

Utilize special protocols and services mated to the cooperative sensing needs. The value of providing specialized capabilities for parallel applications to use has been proven with the collective communications primitives provided in MPI (Message Passing Interface) standard library [7]. Although there will be less regularity in a randomly deployed wireless network than say a hard-wired multiprocessor, one can expect that communications primitives may be developed that achieve significant improvements in efficiency, as well as provide a standard functions for applications developers to draw upon. For example, a “scatter” MPI type of operation among local sensor nodes would benefit from multicasting that incorporates the inherent broadcast nature of the radio medium; similarly, synchronization may be efficiently implemented.

It is noted that this paper has only emphasized nodes cooperating toward the multipoint sensing function, although there are numerous additional functions where cooperation is also appropriate. For example, cooperation is needed for establishing the communications network (topology determination, scheduling, routing, etc.), for localization of the nodes themselves, energy management, data storage and retrieval, and fault management.

4. EXAMPLES

Consider a “Pixellation” class multipoint measurement application, such as the Magnetospheric Constellation Mission. A key phenomenon of interest is characterization of the dynamic magnetopause, or boundary of the Earth’s magnetosphere. Although an oversimplification, a distributed operation might occur as follows. Each nanosatellite can detect with its magnetometer whether or not it is inside or outside the magnetosphere. If each node could communicate with its neighbor, they could determine if they straddle the boundary. Only those nodes that do lie on the boundary are relevant, and all others may drop out of further processing. In essence, an edge detection process

(such as is common in image processing, although with regular pixel tiling) is executed. More generally, contours may be created from gradients between nodes. It is apparent that such processes can be broken into parallel subtasks. Lateral inhibition techniques may be used to exaggerate edge or point effects [8], and “layers” (such as in neural networks) may be formed to capture increasing levels of abstraction.

In the case of a “Beamformation” class application, breaking the processing into subtasks appears more challenging. Often, such processing involves matrix inversions. These may proceed as in conventional parallel processing implementations, such as use of LU factorization and collective communications, but performance benefits are not as obvious. An approach that has been proposed is to split the sensing nodes into say three local clusters, with centralized beamformation in each cluster to produce target bearing. These three bearings (which require small communications) are then combined at the more macro level to localize the target’s position.

5. CONCLUSION

Multipoint measurement systems provide the means to determine characteristics of phenomena that cannot be determined otherwise. Several classes of such systems have been identified, together with examples. Such systems may be developed out of cooperating nodes that exchange information in a wireless network and possibly perform signal combining operations in a distributed processing manner. Creation of intelligent sensor collectives will be enabled through the development of general algorithms for these classes of multipoint measurement applications, including communications primitives that efficiently utilize resources.

REFERENCES

- [1] J. Pottie and William J. Kaiser, “Embedding the internet: wireless integrated network sensors,” *Communications of the ACM*, 43(5), pp. 51-58, May 2000.
- [2] J. R. Agre and L. P. Clare, “Distributed Sensor Networks,” *IEEE Computer*, vol. 33. No. 5, pp. 106-108 (Integrated Engineering column), May 2000.
- [3] NASA – GSFC Magnetospheric Constellation Mission Document, Al Lieberman, NASA/GSFC, STAAC, Code 730, March 30, 1999, <http://stp.gsfc.nasa.gov/missions/mc/magconmissdoc.pdf>.
- [4] http://planetquest.jpl.nasa.gov/TPF/tpf_index.html.
- [5] Joe C. Chen, Kung Yao, and Ralph E. Hudson, “Source Localization and Beamforming,” *IEEE Signal Processing Magazine*, pp. 30-39, March 2002.
- [6] <http://climate.gsfc.nasa.gov/~wiscombe/Leonardo.html>.
- [7] Section on collective communications in <http://www.netlib.org/utk/papers/mpi-book/mpi-book.html>.
- [8] Martin Brooks, “Highly Redundant Sensing in Robotics – Analogies From Biology: Distributed Sensing and Learning,” NATO Advanced Research Workshop on Highly Redundant Sensing in Robotic Systems, Italy 1988.

Loren Clare is the Technical Group Supervisor for the Communications Networks Group at the Jet Propulsion Laboratory. He obtained the Ph.D. in System Science from the University of California, Los Angeles in 1983. His research interests include wireless communications protocols, self-organizing systems, network systems design, modeling and analysis, and realtime distributed control systems. Prior to joining JPL in May 2000, Dr. Clare was a senior research scientist at the Rockwell Science Center. Recent work focused on distributed sensor networks, including low-power protocol design as well as development of a layered architecture for cooperative signal processing. He has extensive experience in tactical network design and analysis, as well as satellite networking. He has also developed communications protocols for realtime networks supporting industrial automation.

